CLAIMS

1. One or more computer readable media having stored thereon a plurality of instructions that, when executed by one or more processors, causes the one or more processors to perform acts including:

receiving an initial image selection;

generating a plurality of query vectors by extracting, for each query vector, one of a plurality of low-level features from the initial image selection;

selecting a set of potentially relevant images based at least in part on distances between the plurality of query vectors and a plurality of feature vectors corresponding to low-level features of a plurality of images;

receiving feedback regarding the relevance of one or more images of the set of potentially relevant images;

generating a new plurality of query vectors based at least in part on the feedback;

generating a weighting of feature elements based at least in part on the feedback; and

selecting a new set of potentially relevant images based at least in part on both the weighting of feature elements and distances between the new plurality of query vectors and the plurality of feature vectors.

2. One or more computer readable media as recited in claim 1, wherein the selecting a new set of potentially relevant images comprises using a matrix in determining the distance between one of the new plurality of query vectors and one of the plurality of feature vectors, and further comprising dynamically selecting the matrix based on both a number of images in the set of potentially relevant images for which relevance feedback was input and a number of feature elements in the one feature vector.

3. One or more computer readable media as recited in claim 2, wherein the dynamically selecting comprises using a diagonal matrix if the number of images in the set of potentially relevant images for which relevance feedback was input is less than the number of feature elements in the one feature vector, and otherwise using a full matrix.

4. One or more computer readable media as recited in claim 2, wherein the dynamically selecting comprises:

if the number of images in the set of potentially relevant images for which relevance feedback was input is not less than the number of feature elements in the one feature vector, then using one matrix that transforms the query vector and the one feature vector to a higher-level feature space and then using another matrix that assigns a weight to each element of the transformed query vector and the transformed feature vector; and

if the number of images in the set of potentially relevant images is less than the number of feature elements in the one feature vector, then using a matrix that assigns a weight to each element of the query vector and the one feature vector.

5. One or more computer readable media as recited in claim 2, wherein X represents an image matrix that is generated by stacking N feature vectors corresponding to the set of potentially relevant images for which relevance feedback was received and resulting in an (N x K) matrix, C represents a weighted covariance matrix of X, det(C) represents the matrix determinant of C, and the matrix comprises a full matrix (W^*) that is generated as follows:

$$W^* = (\det(C))^{\frac{1}{K}} C^{-1}.$$

6. One or more computer readable media as recited in claim 2, wherein w_{kk} represents the kk^{th} element of matrix W, x_k represents the k^{th} feature element, σ_k represents the standard deviation of the sequence of x_k's, the matrix comprises a diagonal matrix with each diagonal element (w_{kk}) being generated as follows:

$$w_{kk} = \frac{1}{\sigma_k}$$

7. One or more computer readable media as recited in claim 1, wherein N represents the number of images in the set of potentially relevant images for which relevance feedback has been received, π_n represents the relevance of image n in the set of images, $\vec{\pi}^T$ represents a transposition of a vector generated by concatenating the individual π_n values, and X represents an image matrix that is generated by stacking N training vectors corresponding to the set of potentially relevant images into a matrix, and wherein each new query vector ($\vec{q}$) of the new plurality of query vectors is generated as follows:

$$\vec{q} = \frac{\vec{\pi}^T X}{\sum_{n=1}^{N} \pi_n}.$$

8. One or more computer readable media as recited in claim 1, wherein f_i represents a summation, over the images in the set of potentially relevant images, of a product of a relevance of the image and a distance between the query vector and the feature vector, and wherein the selecting a new set of potentially relevant images comprises combining, for each image, a weighted distance between the plurality of query vectors and the plurality of feature vectors, and wherein the weight (u_i) for each of a plurality (I) of distances between a query vector and a corresponding feature vector is calculated as:

$$u_i = \sum_{j=1}^{I} \sqrt{\frac{f_j}{f_i}}.$$

9. One or more computer readable media as recited in claim 1, wherein the receiving feedback comprises receiving feedback from a user.

10. One or more computer readable media as recited in claim 1, wherein the low-level features include: a color moments feature, a wavelet based texture feature, and a water-fill edge feature.

11. A method of selecting between two types of matrixes to be used to weight, based on relevance feedback, a plurality of feature elements for image retrieval, the method comprising:

selecting one of the two types of matrixes based on both a number of previously retrieved relevant images and a length of a feature vector including the plurality of feature elements.

12. A method as recited in claim 11, wherein the selecting comprises selecting one of the two types of matrixes based on both a number of previously retrieved potentially relevant images which were identified by a user as being relevant, and the length of the feature vector including the plurality of feature elements.

13. A method as recited in claim 11, wherein the plurality of feature elements are all elements of the same feature.

14. A method as recited in claim 11, wherein the selecting comprises using a first type of matrix if the number of retrieved relevant images is less than the length of the feature vector, and otherwise using a second type of matrix.

15. A method as recited in claim 14, wherein the first type of matrix comprises a diagonal matrix and wherein the second type of matrix comprises a full matrix.

16. A method as recited in claim 11, wherein the selecting comprises using a first type of matrix if the length of the feature vector exceeds the number of retrieved relevant images by at least a threshold amount, and otherwise using a second type of matrix.

17. A method as recited in claim 16, wherein the first type of matrix comprises a full matrix and the second type of matrix comprises a diagonal matrix.

18. One or more computer readable media including a computer program that is executable by a processor to perform the method recited in claim 11.

19. One or more computer readable media having stored thereon a plurality of instructions that, when executed by one or more processors, causes the one or more processors to perform acts including:

comparing, for each of a plurality of images, a plurality of feature elements from a query vector to a plurality of feature elements from a feature vector corresponding to the image;

identifying a number of potentially relevant images based on the comparing;

receiving user feedback regarding the relevancy of one or more of the potentially relevant images;

re-comparing, for each of the plurality of images, the plurality of feature elements from the query vector to the plurality of feature elements from the feature vector, including using a matrix to compare the feature elements and dynamically selecting a type of matrix to use based on both the user feedback and the number of the plurality of feature elements;

identifying a new set of potentially relevant images based on the re-comparing; and

presenting the new set of potentially relevant images to the user.

20. One or more computer readable media as recited in claim 19, wherein the re-comparing comprises dynamically selecting the type of matrix to use based on both a number of the potentially relevant images for which user feedback has been received and the number of the plurality of feature elements.

21. One or more computer readable media as recited in claim 19, wherein the dynamically weighting comprises using a first type of matrix if the number of retrieved relevant images is less than the length of the feature vector, and otherwise using a second type of matrix.

22. One or more computer readable media as recited in 21, wherein the first type of matrix comprises a diagonal matrix and the second type of matrix comprises a full matrix.

23. A method comprising:

generating a query vector corresponding to a feature of one image;

identifying a feature vector corresponding to the feature of another image;

identifying a number of training samples for which relevance feedback has been received;

if the number of training samples either equals or exceeds a threshold amount, then determining a distance between the query vector and the feature vector including transforming the query vector and the feature vector to a higher-level feature space and then assigning a weight to each element of the transformed query vector and the transformed feature vector; and

if the number of training samples does not exceed the threshold amount, then determining the distance between the query vector and the feature vector including assigning a weight to each element of the query vector and the feature vector.

24. A method as recited in claim 23, wherein the feature vector includes a plurality of feature elements and wherein the threshold amount comprises the number of feature elements in the feature vector.

25. A method as recited in claim 23, wherein if the number of training samples either equals or exceeds the threshold amount, then determining the distance (g), where P is a mapping matrix, $\vec{q}$ is the query vector, $\vec{x}$ is the feature vector, and Λ is a weighting matrix, as:

$$g = (P(\vec{q} - \vec{x}))^T \Lambda (P(\vec{q} - \vec{x})).$$

26. A method as recited in claim 23, wherein if the number of training samples does not exceed the threshold amount, then determining the distance (g), where $\vec{q}$ is the query vector, $\vec{x}$ is the feature vector, and Λ is a weighting matrix, as:

$$g = (\vec{q} - \vec{x})^T \Lambda (\vec{q} - \vec{x}).$$

27. A method as recited in claim 23, further comprising:

repeating the generating, identifying of the feature vector, identifying of the number of training samples, and the determining for each of a plurality of features; and

identifying how closely the image and the other image match each other by combining the distances between the query vectors and the feature vectors for the plurality of features.

28. A method as recited in claim 27, wherein the identifying comprises calculating a weighted summation of each of the individual distances for each of the plurality of features.

29. One or more computer readable media including a computer program that is executable by a processor to perform the method recited in claim 23.

30. A system comprising:

a query vector generator to generate a query vector corresponding to a feature of one image;

a comparator, coupled to the query vector generator, to,

identify a feature vector corresponding to the feature of another image,

identify a number of training samples for which relevance feedback has been received,

if the number of training samples either equals or exceeds a threshold amount, then to determine a distance between the query vector and the feature vector including transforming the query vector and the feature vector to a higher-level feature space and then assigning a weight to each element of the transformed query vector and the transformed feature vector, and

if the number of training samples does not exceed the threshold amount, then to determine the distance between the query vector and the feature vector including assigning a weight to each element of the query vector and the feature vector.

31. A method comprising:

for one of a plurality of images and each of a plurality of features,

generating, based on the set of search criteria, a query vector for the feature,

identifying a feature vector, corresponding to the image, for the feature, and

determining how closely the feature vector matches the query vector;

determining how closely the image matches the set of search criteria based on how closely, for the plurality of features, the feature vectors match the query vectors.

32. A method as recited in claim 31, wherein generating the query vector comprises generating the query vector based at least in part on user relevance feedback regarding how relevant images previously displayed to a user were.

33. A method as recited in claim 31, wherein identifying the feature vector comprises:

identifying a low-level feature vector corresponding to the feature; and

mapping the low-level feature vector to a higher level feature space.

34. A method as recited in claim 33, wherein the identifying the feature vector further comprises incorporating, into the mapping, relevance feedback.

35. A method as recited in claim 31, wherein the initial search criteria comprises an image.

36. A method as recited in claim 31, wherein the determining how closely the feature vector matches the query vector comprises determining a distance between the feature vector and the query vector, and wherein the determining how closely the image matches the set of search criteria comprises calculating a weighted summation of each of the individual distances between the feature vectors and the query vectors.

37. A method as recited in claim 36, wherein the calculating a weighted summation comprises calculating the weighted summation based at least in part on user relevance feedback regarding how relevant images previously displayed to a user were.

38. One or more computer readable media including a computer program that is executable by a processor to perform the method recited in claim 31.

39. One or more computer readable media having stored thereon a plurality of instructions that, when executed by one or more processors, causes the one or more processors to perform acts including:

identifying a plurality of query vectors for one image, each query vector corresponding to one of a plurality of features;

identifying a plurality of feature vectors for the other image, each feature vector corresponding to one of the plurality of features;

for each feature, determining a distance between the corresponding query vector and the corresponding feature vector; and

combining the distances to generate a value representing an overall distance between the two images.

40. One or more computer readable media as recited in claim 39, wherein the identifying the plurality of query vectors comprises extracting the plurality of query vectors from the image.

41. One or more computer readable media as recited in claim 39, wherein the identifying the plurality of query vectors comprises generating the plurality of query vectors based at least in part on user relevance feedback regarding how relevant images previously displayed to a user were.

42. One or more computer readable media as recited in claim 39, wherein the determining the distance between the corresponding query vector and the corresponding feature vector includes incorporating, into the determining, user relevance feedback regarding how relevant images previously displayed to a user were.

43. One or more computer readable media as recited in claim 39, wherein the combining the distances comprises calculating a weighted summation of each of the individual distances between the feature vectors and the query vectors.

44. One or more computer readable media as recited in claim 43, wherein the calculating a weighted summation comprises calculating the weighted summation based at least in part on user relevance feedback regarding how relevant images previously displayed to a user were.

45. A system comprising:

a query vector generator to identify a query vector, for each of a plurality of features, corresponding to one image;

a comparator, coupled to the query vector generator, to,

identify, for each of a plurality of features of another image, a feature vector for the feature,

determine, for each of the plurality of features, how closely the feature vector matches the query vector, and

determine how closely the image matches the other based on how closely, for the plurality of features, the feature vectors match the query vectors.

46. A method of generating a query vector to compare to a feature vector of another image, the method comprising:

receiving feedback regarding the relevance of each image of a set of images;

wherein N represents the number of images in the set of images for which user relevance feedback has been received, π_n represents the relevance of image n in the set of images, $\vec{\pi}^T$ represents a transposition of a vector generated by concatenating the individual π_n values, and X represents an image matrix that is

generated by stacking N training vectors corresponding to the set of images into a matrix; and

generating a query vector ($\vec{q}$) corresponding to one of a plurality of features as follows:

$$\vec{q} = \frac{\vec{\pi}^{T} X}{\sum_{n=1}^{N} \pi_n}.$$

47. One or more computer readable media including a computer program that is executable by a processor to perform the method recited in claim 46.

48. A method of generating a weight to apply to distances between query vectors and feature vectors when combining the distances, the method comprising:

receiving feedback regarding the relevance of each image of a set of images;

wherein f_i represents a summation, over the images in the set of images, of a product of a relevance of the image and a distance between the query vector and the feature vector; and

generating a weight (u_i) for each of a plurality (I) of distances between a query vector corresponding to one of a plurality (I) of features and a feature vector corresponding to the one of the plurality (I) of features as:

$$u_i = \sum_{j=1}^{I} \sqrt{\frac{f_j}{f_i}}.$$

49. One or more computer readable media including a computer program that is executable by a processor to perform the method recited in claim 48.

50. A system comprising:

a client device;

a collection of a plurality of images;

an image server, coupled to the client device and the collection of a plurality of images, the image server to receive image retrieval requests from the client device and to,

receive an initial image selection from the client device,

generate a plurality of query vectors by extracting, for each query vector, one of a plurality of low-level features from the initial image selection,

select a set of potentially relevant images based at least in part on distances between the plurality of query vectors and a plurality of feature vectors corresponding to low-level features of a plurality of images,

receive feedback regarding the relevance of one or more images of the set of potentially relevant images,

generate a new plurality of query vectors based at least in part on the feedback,

generate a weighting of feature elements based at least in part on the feedback, and

select a new set of potentially relevant images based at least in part on both the weighting of feature elements and distances between the new plurality of query vectors and the plurality of feature vectors.

Add a1

Add b1